1

DESCRIPTIVE, NORMATIVE, AND PRESCRIPTIVE INTERACTIONS IN DECISION MAKING

DAVID E. BELL, HOWARD RAIFFA, AND AMOS TVERSKY

The focus of our attention is the individual decision maker facing a choice involving uncertainty about outcomes. We will consider how people do make decisions, how "rational" people should make decisions, and how we might help less rational people, who nevertheless aspire to rationality, to do better. When we speak of nonrational people, we do not mean those with diminished capacities; we refer instead to normal people who have not given thought to the process of decision making or, even if they have, are unable, cognitively, to implement the desired process. Our decision makers are not economic automatons; they make mistakes, have remorse, suffer anxieties, and cannot make up their minds. We start with a premise, not that people have well thought out preferences, but that they may be viewed as having divided minds with different aspirations, that decision making, even for the individual, is an act of compromise among the different selves.

For our purposes we shall augment the usual dichotomy that distinguishes between the normative and descriptive sides (the "ought" and the "is") of decision making, by adding a third component: the *prescriptive* side. We do this because much of our concern in this paper addresses the question: "How can real people – as opposed to imaginary, idealized, super-rational people without psyches – make better choices in a way that does not do violence to their deep cognitive concerns?" And we find that much that we have to say on these matters does not fit conveniently into the usual normative or descriptive niches. Loosely speaking, prescriptive analyses exploit some of the logical consequences of normative theories and the empirical findings of descriptive studies but, in addition, something else has to be added that is far from the spirit of normative or descriptive analyses.

In order to contrast the three modes of analysis let us consider the concept of transitivity. Insofar as decision makers have opinions about alternatives most (but not all) normative models posit that if the decision agent prefers alternative A to B, and prefers B to C, then he will also have a preference for A over C. This is a common axiom or desideratum of many normative systems. When the

transitivity axiom is added to other axioms in a given normative framework, the analyst can generate a superstructure of logical implications. This is the familiar world of the mathematician. Axioms are not God-given but are chosen by the mathematical creator who has one eye on the real world for inspiration. In the case of transitivity the plausibility is self-evident; people, by and large and most of the time, make choices in a transitive fashion.

But describers of reality point out that most individuals occasionally exhibit intransitive (or cyclical) preferences: A over B, B over C, and C over A. There are lots of ways to explain why this might be the case. Preferences might be nonstationary; or they might be stochastic; or they might involve a balancing of attributes. When A and B are compared, the decision maker might highlight certain underlying attributes. Then when the choice set changes from (A, B) to (A, C) or to (B, C) different attributes might be emphasized. Or perhaps the decision maker has consulted three experts in order to make up his own mind and the experts have exhibited the classical cyclical pattern: A over B over C for expert 1, B over C over A for expert 2, and C over A over B for expert 3. When the decision maker combines his advice using a majority rule principle, the intransitivity becomes manifest: A over B, B over C, and C over A. Or perhaps there are no external experts but within the psyche of the single decision maker we can imagine that there are three selves (self′, self″, self‴) whose aggregate preferences are cyclic. The point of all this is that, descriptively speaking, people are often intransitive and there is a divergence between those normative conceptions of choice that posit transitivity and observed behavior. The normatively oriented analyst can try to accommodate some degree of intransitivity in his mathematical, axiomatic abstractions and attempt to fuse a better concordance between idealistic theorizing and empirical behavior. But perhaps even though some individuals exhibit intransitivities they wish to act otherwise. Is this the case? Well, this again is an empirical proposition that psychologically minded decision theorists may wish to investigate.

Now let us add the prescriptive viewpoint. Mr Johnson has racked his mind and has eliminated a lot of alternatives but is now perplexed about whether he should choose alternative A or C. Choose he must. The trouble is that the alternatives are so different in so many different dimensions and, in addition, uncertainties complicate the picture. Now let us suppose that Johnson (or his consulting decision analyst), after examining aspects of A and C, ingeniously invents a new alternative B for which Johnson finds it comfortable to say that he prefers A to B and B to C. Johnson might also think it is reasonable that his preference involving A, B, and C should be transitive. So the creation of hypothetical alternative B might help Johnson to believe that deep down he really prefers A to C. Alternatively, instead of the introduction of B, there might be the introduction of B', B'', and B''' such that $A > B'$, $B' > B''$, $B'' > B'''$, and $B''' > C$, where $>$ means "is preferred to."

Is this decision-aiding device descriptive? If it were, Johnson would do this for himself. Is it normative (A is preferred to C if and only if there exists B, etc.)? No,

we do not think so. What the decision analyst did was find a way to help Johnson by calling on a normative device, in this case transitivity. A prescriptive analyst should also be sophisticated enough to know that Johnson might not really prefer A to C deep down but he might have been led by magic mirrors into believing that he does. Can the same trick be played in reverse? Can an alternative D be concocted such that $C > D$, and $D > A$, which, by transitivity, would lead to the conclusion that Johnson really prefers C to A? Well, it is worth a try. This is the spirit of the art of prescriptive analysis; it is a curious mixture of normative and descriptive analysis but a lot more.

Another example will elucidate the distinctions between the normative, descriptive, and prescriptive viewpoints. McNeil *et al.* (chapter 26 in this volume)

> investigated how variations in the way information is presented to patients influence their choices between alternative therapies. Data were presented summarizing the results of surgery and radiation therapy for lung cancer to 238 ambulatory patients with different chronic medical conditions and to 491 graduate students and 424 physicians.

The subjects were asked to imagine that they had lung cancer and had to choose between two therapies on the basis of given probabilistic assessments. The relevant question (repeated below) was presented to one subset of individuals in terms of survival.

> Of 100 people having treatment A, 90 live through the treatment. A total of 70 people are alive by the end of the first year and a total of 38 people are alive by the end of five years.
>
> Of 100 people having treatment B, all live through the treatment. A total of 79 people are alive by the end of the first year and a total of 26 people are alive by the end of five years.

A second group of individuals was presented with the same choice, except that the data were described in terms of mortality rates rather than survival rates (e.g., "of 100 people having treatment A, 10 die during treatment," etc.).

The data, using the appropriate frame of survival or mortality, are also tabulated and shown to each subject, who is then asked to choose treatment A or B (see table 1.1, frames 1.1 and 1.2). Each subject is shown either frame 1.1 or 1.2 but not both. Observe that the data given to the subjects are identical in informational content whether they are given in terms of survival or mortality.

The punch line is dramatic. There are vast differences in the responses of subjects depending on whether the data are presented in terms of survival or in terms of mortality. The format seems to influence the thought process. The proportion of subjects that preferred treatment B to A was 61% when the data were presented in terms of mortality and was only 37% when presented in terms of survivability. These are the descriptive realities. Choices can be influenced by the framing of the question, a major concern that is not incorporated in the usual normative theory, but is well known among students of the field.

Prescribers, allegedly helpful intervenors in the decision-making process, have

Table 1.1. *Alternate frames for a choice selection*

	(Frame 1.1) Cumulative chance of death		(Frame 1.2) Cumulative chance of survival	
	Treatment *A*	Treatment *B*	Treatment *A*	Treatment *B*
During treatment	10%	0%	90%	100%
By year 1	30%	21%	70%	79%
By year 5	62%	74%	38%	26%

Outcome		(Frame 1.3) Treatment *A*	Treatment *B*
(0)	Not surviving treatment (i.e., dying during treatment)	10%	0%
(0, 1)	Surviving treatment but dying before end of first year	20%	21%
(1, 5)	Surviving 1 year but dying before end of fifth year	32%	53%
(5+)	Surviving at least 5 years	38%	26%
		100%	100%

to be continually aware that they can bias the responses of subjects by the mere wording of questions. In the above example, one possibility is to join the two frames together in the description of a consequence (e.g., "90 live through the treatment and therefore 10 die"). McNeil, Pauker, and Tversky (chapter 26 in this volume) presented some subjects with both versions of the questionnaire and examined subjects' responses to the combined frame and their reactions to their own inconsistencies. The effect of the combined frame was much closer to the mortality frame than to the survival frame, perhaps because the mortality data are more salient than the survival data. When presented with their inconsistencies most subjects are inclined to modify their choices, but are not clear as to which preference should be changed. These observations are descriptively interesting and prescriptively important if we are going to give advice to real people, but they are hardly relevant in most normative conceptions.

Let us continue with this illustration from a prescriptive orientation. The data can also be summarized in a composite version in table 1.1 (frame 1.3) that gives the relative frequency (chance) of each outcome.

Now let us use some more prescriptive magic on a hypothetical decision maker Ms Jones who is faced with frame 1.3. Imagine that we have an urn with 100 balls, each ball having an *A* label on it and a *B* label on it. The *A* labels may be marked in red and the *B* labels in blue. In table 1.2 we indicate possible markings on the 100 balls that correspond to the frequencies in table 1.1. For example, 20 balls

Table 1.2

Number of balls	*A*-label	*B*-label
20	(0, 1)	(0, 1)
32	(1, 5)	(1, 5)
26	(5+)	(5+)
1	(0)	(0, 1)
9	(0)	(1, 5)
12	(5+)	(1, 5)

Table 1.3

	Original Treatments		Reduction Treatments		Modification Treatments	
Event	*A*	*B*	*A'*	*B'*	*A''*	*B''*
(0)	10	0	10	0	11	0
(0, 1)	20	21	0	1	0	0
(1, 5)	32	53	0	21	0	22
(5+)	38	26	12	0	11	0
	100	100	22	22	22	22

have a common A and B label, namely (0, 1); and in the last row, 12 balls have an A label of (5+) and a B label of (1, 5). The problem can be viewed as follows: Ms Jones must make a choice of A or B and then draw a ball at random that will determine the outcome.

But there are 78 balls that have identical A and B labels and these balls, so it can be argued, are just cluttering up the problem. So let us get rid of them in order to concentrate on the essence of the choice problem: the choice between the labels on the remaining 22 balls. The original and reduced problems are shown in table 1.3. The argument goes that the choice between A and B should be the same as the choice between A' and B'. Of course, by and large people might choose differently between these two versions, but should they? Would you? Suppose Ms Jones says on reflection that she would want her choice between A and B to be governed by her choice between A' and B' and that she would like to think hard about her choice in the latter problem. Furthermore, let us assume that she is tilting toward favoring A' over B', but is not sure. To test her preference let us pose the following question to Ms Jones. Let us make A' a bit worse to yield A'' and improve B' a bit to yield B'' (see the last two columns of table 1.3). Now let us ask Ms Jones: "Suppose you are told that you can be certain of the outcome (1, 5); would you take a 50–50 chance of getting (0) or (5+)?"

Ms Jones thinks: "Sure, I would do it." Therefore she prefers A'' to B'' and therefore she *should* prefer A' to B' and therefore she *should* prefer A to B. There

are a lot of "shoulds" lined up. Is this bit of trickery of help to Jones? Well, it might help us to think deeper about the problem but Jones might think otherwise. She might ask, "If I announce A and pick a ball would I be told the outcome on the B-label as well?" After all the comparison of what she got and could have got might have an effect on her psyche. Should it? Well even if it *should not* in your opinion, what happens if it *does* in her opinion? This, and a lot more, is what the prescriptive intervenor must wrestle with.

When problems are recast in supposedly equivalent forms there may be subtle psychological nuances that are omitted that on cognitive reflection should not be omitted. Prescriptive intervenors should be especially sensitive to these nuances and they can be sensitized somewhat to understand these nuances by comparing descriptive and normative behaviors. It is much more difficult, however, to know what to do about these nuances in prescriptive interventions.

It is our aim in this chapter to highlight discrepancies between real and idealized behavior such as previewed in the above two motivating examples and to throw out the challenge: what should be done to improve behavior of real people? That is not the aim of those who think of themselves as normative modelers, nor the aim of most empiricists who investigate real-world behavior. Decision analysts who are interested in prescriptive interventions must rise to the challenge if on balance they are to do more good than harm.

Before we attempt a little more systematically to describe potential interactions among descriptive, normative, and prescriptive perspectives, with the purpose of highlighting an identifiable niche for the prescriptive category, we will in the next section put aside questions of prescription and discuss the very fuzzy line separating descriptive and normative analyses.

NORMATIVE IDEALIZATIONS USED AS DESCRIPTIVE APPROXIMATIONS

Mathematical economics – or should we say simply "economics" because so much of economics is mathematical these days – makes extensive use of models that posit utility maximization behavior on the parts of individual agents. A typical example is of the form: let $\underline{x}_i = (x_{1i}, \ldots, x_{ji}, \ldots, x_{Ji})$ represent the allocation of J resources enjoyed by individual i. Let $u_i(\underline{x}_i)$ be i's utility evaluation for $\underline{x}_i$, where u_i is assumed to be strictly monotonic in each variable, concave, and twice differentiable; let . . . etc. Then the abstraction continues by letting the I individuals of this "economy" interact, i.e., trade commodities and information; it posits that each agent seeks to maximize his own expected utility given some rational expectations of what others might do. Is this abstraction normative or descriptive? Or perhaps neither? For the most part researchers who write in this genre are not doing it to help guide or prescribe behavior for agent i. Rather such models are usually formulated as first-cut approximations of the descriptive behavior of individuals. Even though the vast majority, if not all, of the subjects

do not behave coherently enough to satisfy the implied behavior of the assumptions of the model, the feeling is that the model might yield empirically meaningful insights because if real-world behavior deviated too far from the equilibrium behavior suggested by the model, a few individuals would exploit such aberrant behavior to their own advantage, and through learning and adaptation the economy would settle down to some distinguishable equilibrium state predicted by the model. In this manner, the normative character of the model is used as an argument to reinforce its descriptive value. A second-cut descriptive model, a far more difficult modeling exercise, would be to posit quasi-rational motives for individual actors and make the model dynamic with adaptive learning feedbacks. A third-cut descriptive model might try to better understand a much richer panoply of human motivations and emotions so that what is deemed "quasi-rational" behavior might indeed be more fully rationalized.

The theoretical results emanating from the simpler model that posits the existence of utility functions only over physical commodity bundles could be more robust than the assumptions themselves. Much of this modeling effort is descriptive – not descriptively accurate in the small but descriptively suggestive or approximative in the large. More realistic and descriptively plausible micro-behavior can be grafted onto the model, although this is not commonly done.

Some researchers who design and analyze such abstract models of economies with idealized agents might have a macro-prescriptive motivation in mind, such as: how should information flows among the agents be modified to make this small economy perform more efficiently? But very often that prescriptive orientation is left unexpressed in the model. The supra-decision maker, the rules manipulator, is not introduced in the model, not given a utility function over the multi-attributed societal concerns; the model is used in its formal structure as an approximation of descriptive reality and is used in a casual, informal manner for a diffuse macro-prescriptive purpose.

Economics textbooks, for pedagogical purposes, assume that economic agents have indifference curves in n-space. These indifference maps are just assumed to exist and no serious attempt is made at establishing a protocol for eliciting such indifference maps as would be necessary if the purpose of the analysis were really prescriptive – prescriptive in the sense of trying to guide reflective behavior. The purpose of the textbook writers really is first-cut descriptive behavior and not prescriptive behavior. For some purposes it may be helpful to foster the myth that deep down in all of us there is an orderly, coherent, n-dimensional preference function and that, when we elicit manifest, phenotypic responses from this latent, genotypic preference function, errors are made. That myth, however, may not be helpful for prescribing (guiding) behavior of the individual. We believe that such orderly representations (such as n-dimensional utility functions and multivariate subjective probability distributions) do not really exist but in some circumstances their existence might be usefully fabricated and constructively developed in order to guide behavior.

DESCRIPTIVE, PRESCRIPTIVE, AND NORMATIVE ANALYSES

We have refrained up to now from defining just what we mean by descriptive, prescriptive, and normative analyses because it is a bit easier for us to back into a clarification of the distinctions we wish to make. Some authors use "normative" and "prescriptive" interchangeably but we propose to treat these terms quite distinctly.* Our emphasis in this chapter will largely be on prescriptive analysis and therefore it is important for us to distinguish it clearly from normative analysis.

Let us start with *descriptive* analysis because this is easiest. How do real people think and behave? How do they perceive uncertainties, accumulate evidence, learn and update perceptions? How do they learn and adapt their behavior? What are their hang-ups, biases, internal conflicts? How do they talk about their perceptions and choices? Do they really do as they say they do? Can they articulate the reasons for their actions? How do they resolve their internal conflicts or avoid such resolutions? Do they decompose complex problems, think separately about component parts of problems, and then recompose or integrate these separate analyses? Or do they think more holistically and intuitively? What are the differences in types of thought patterns for people of different cultures, of different experience levels? What is the role of tradition, imitation, superstition in decision making (or nonmaking)? How can "approximate" real behavior be described? How good are various mathematical models in predicting future behavior?

In short, descriptive analysis is concerned with how and why people think and act the way they do. At times it may involve intricate mathematical modeling and require sophisticated statistical analysis. It is a highly empirical and clinical activity that falls squarely in the province of the social sciences concerned with individual behavior. Scholars can study this domain without any concern whatsoever of trying to modify behavior, influence behavior, or moralize about such behavior.

Now for the normative side. This activity is harder to characterize because it involves several facets. First, there is the notion that normative theory has something to do with how idealized, rational, super-intelligent people should think and should act. Such analyses abstract away known cognitive concerns of real people, their internal turmoils, their shifting values, their anxieties and lingering post-decisional disappointments and regrets, their repugnance (or zest) for ambiguity or danger, their inabilities to do intricate calculations, and their limited attention span. The hallmarks of such normative analyses are coherence and rationality as captured usually in terms of precisely specified desiderata or axioms of the form: if the decision maker believes so and so, he should do such and such. As usual in any mathematical system, the power of any set of desiderata comes from their logical, synergistic, joint implications.

Axioms, basic principles, and fundamental desiderata are motivated by what

* For the purposes of this chapter only; we are not attempting to change common usage.

some investigator thinks is logical, rational, intelligent behavior. Then like any mathematical axiom system (such as sets of axioms for geometry) the academic researchers play variations on the themes: what happens if this axiom is dropped, or if this axiom is modified in such and such a way? This exercise is rewarding if the mathematical implications are profound or aesthetically pleasing. The exercise can also be rewarding if the researcher can see a better concordance between the abstract system and some aspects of behavior that is empirically verifiable or that the researcher imagines is verifiable. Thus there is a dynamic interaction between the real world, imaginations about the real world, and the abstract mathematical system. There are extant a host of abstract models of decision making bearing some relation to decision making as it is, or as it is perceived to be, or as it should be in someone's mind.

In the usual parlance, an abstract system that purports to describe or predict behavior is called a descriptive model; an abstract system that attempts to capture how ideal people might behave is called a normative model. There is little difficulty in categorizing some models as clearly descriptive or normative. One trouble is that some normatively motivated models are often used, as mentioned above, as first-cut descriptive models. Other clearly normatively motivated models go through successive modifications that try to make them more useful for descriptive and predictive purposes and then it may be difficult to say whether these modifications should be classified as normative or descriptive. On the other hand, some descriptively motivated models are occasionally modified to come a bit closer to what some analyst believes is a proper norm for wise behavior. And then the model falls into the grey area. Is it descriptive or is it normative?

Now we move on to prescription. What should an individual do to make better choices? What modes of thought, decision aids, conceptual schemes are useful – useful not for idealized, mythical, de-psychologized automata – but for real people? And since real people are different, with differing psyches and emotions, capabilities, and needs, good advice has to be tuned to the needs, capabilities, and emotional makeups of the individuals for whom the prescriptive advice is intended. It becomes even more complicated when individuals who think one way have to interact with experts who think along different paradigmatic lines, as, for example, between a rational decomposer and a holistic intuiter.

For some individuals a wise prescriptive might be: "Behave as you normally do. You're doing well and any new mode of analysis might inhibit your creative thinking." For others the advice might be: "It's important that you decompose your problem and get external advice from experts on such-and-such a component part, because otherwise you will not be able to constructively integrate and synthesize what you know together with what others know."

The differences among the three functions – descriptive, normative, and prescriptive – of choice models can be illuminated by examining the criteria by which they are evaluated. Descriptive models are evaluated by their *empirical validity*, that is, the extent to which they correspond to observed choices.

Normative models are evaluated by their *theoretical adequacy*, that is, the degree to which they provide acceptable idealizations or rational choice. Prescriptive models are evaluated by their *pragmatic value*, that is, by their ability to help people make better decisions. To be sure, all three criteria are difficult to define and evaluate, as any student of the philosophy of science knows too well. It is evident, nevertheless, that the criteria are different; an argument against a normative model need not be an argument against a descriptive model and vice versa.

For example, consider the property of stochastic dominance. Because this condition is regarded as a cornerstone of rational choice, any theory that does not obey it can be regarded as unsatisfactory from a normative standpoint. A descriptive theory, on the other hand, is expected to account for observed violations of stochastic dominance (e.g., problems 2 and 8 in Tversky and Kahneman, 1986). A prescriptive analysis may develop procedures designed to eliminate and reduce such violations. A failure of dominance, therefore, can serve as a counter-example to a normative model, as an observation to be explained by a descriptive model, and as a challenge for a prescriptive model.

CONCEPTIONS OF CHOICE

James March (1978) questions the usefulness of choice as a pervasive metaphor for describing and interpreting human behavior and he questions what we shall call the *canonical paradigm* of decision making that posits:

an identified decision agent;
a prespecification of alternative choices in the purview of the decision agent;
a set of potential consequences that can be anticipated and evaluated (ranked) in terms of stable, well-defined objectives;
a partition of the possible states of the world – an articulation of mutually exclusive, collectively exhaustible, possible resolutions of uncertainty with no unanticipated surprises;
information and evidence that can be accumulated for the relevance it has for the choice process.

March observes that the way decisions are talked about is not necessarily the way decisions are made. He asserts that, "Our theoretical ideas about choice are partly inconsistent with what we know about human processes of decision and that as a result we sometimes fail to understand what is going on in decision making, *and that as a consequence we sometimes offer less than perfect counsel to decision makers*" (italics added).

We certainly concur with the sentiments he expresses in *describing* what is happening out there, even though we may not know what to do about it; we also believe that March's observations are pertinent to prescribers or intervenors who want to influence the way decisions might be made better or more wisely.

What we call the "canonical paradigm of choice behavior" – an identified decision maker with prespecified alternatives, consequences, states-of-the-world,

preferences, and beliefs – March refers to as a "willful-choice model." He observes and bemoans the fact that such models are omnipresent not only in modern economics but in large parts of anthropology, psychology, political science, and sociology as well as the applied fields that build upon them.

Willful-choice models may not answer questions about what happened in an organization or in society – "Why did it happen?" or "Why did you do it?" – but they might address the question posed by some perplexed actor of that society at a critical juncture in his or her life, "What should I do next?" Societal outcomes result usually from a concatenation of actions, a few taken willfully but others by tradition, by obligation, by duty, by inaction or default, and some taken by chance or God. A single individual rarely makes a grandiose choice for society. But all of us make less grandiose micro-choices: to buy or not to buy, to vote for X or Y, to take job Q or S, and so on. Perhaps each of us, over the course of a lifetime, make a dozen or so critical, deliberative, serious choices. That would add to a vast pool of potential applications for the prescriptive uses of the willful-choice paradigm. But, granted, it would still be a terrible distortion to describe the dénouement of most lifetime careers in terms of these deliberative choices. A lot else happens that can be better described in other ways.

It is alleged that willful-choice models posit the prior specification of *all* alternatives and only rarely does one have all the alternatives. Well, that is not such a serious problem. It still may be helpful to choose as wisely as we can among just those two or three alternatives we can think of *a priori* even though there may be other better alternatives that we are not wise enough to think about. A much more serious objection may be that the decision agent may concentrate so much time on choosing among preconceived alternatives that not enough time or effort goes towards devising new alternatives. That may be the case, but formalizing that meta-dilemma may not be prescriptively productive. It also should be noted, however, that a systematic, deliberative analysis of a set of preconceived alternatives may itself spark creative insights that could generate other more imaginative alternatives.

Herbert Simon and others stress that complex problem solving involves a search process and, descriptively speaking, intelligent decision agents adaptively set up aspirations for this search process. They satisfice. Descriptively this may be the way most individuals behave. Normative models might attempt to rationalize the search-and-quitting process in terms of subjective expected utility maximization in which time and physical and emotional effort are included as components of the utility calculus. This is an old debate between students with normative and descriptive persuasions. This problem is, however, critical for the prescribers: how best to give advice about the search process. A prescriber might believe that satisficing behavior could be adequately rationalized by the maximization of a suitably complex objective function but, if it is too horrendously complicated to constructively formulate such an objective function, then the prescriber's operational advice might be better organized by a satisficing heuristic.

In willful-choice models, preferences for consequences are usually posited to be stable, consistent, precise, and exogenous to the problem. The truth of the matter is that they are often ill-formed, labile, shifting, and endogenous to the problem. Anecdotes galore can be cited where a decision agent feels for some amorphous, nonarticulatable reason that decision A is better than B. It just feels better. His articulated preferences for consequences are not antecedent to his action but rather are derivatives from his action preference. But this is not always the case, and when it is not the case a decision maker might want to think systematically about his basic preferences and base his choice on the implications of these deliberations. Keeney and Raiffa (1976) cite an example in which a decision maker, having already decided what action was best, nevertheless chose to investigate systematically his values and beliefs in order to help develop an advocacy document for his preferred alternative. Surprisingly, an analysis that systematically probed his preferences shifted his opinion. There is no shortage of anecdotes on both sides of this debate.

It is not our aim in this chapter to make an exhaustive list of all the objections to the canonical willful-choice model. We want to underscore, however, our belief that most objections to the use of willful-choice models for the selection of actions yet to be made pose deep intellectual challenges that fall more in the domain of prescriptive analysis than in normative analysis.

THE SUBJECTIVE EXPECTED UTILITY (SEU) MODEL OF CHOICE BEHAVIOR

We now turn our attention to one of several models of willful-choice behavior, the one that is most extensively applied and most often maligned: the subjective expected utility (SEU) model. It is often the case that specific criticisms leveled against the SEU model should more appropriately be directed against the broader class of willful-choice models.

In the SEU model it is assumed that any action chosen by the decision maker will result in some consequences whose specification involves no uncertainties – all the uncertainties of the problem are loaded into what is generally termed "states" or "states-of-the-world." Thus the decision maker confronts an array of states-of-the-world, one of which will ultimately prevail and, given his usually vague information about which of these states will prevail, he must choose an action. It is assumed that the action he chooses will yield some consequence depending on the state that providence, so to speak, selects. In looser parlance, the choice of an act results in a lottery that will yield one consequence depending on which state prevails. The decision maker thus has a choice over lotteries. Consequences and states-of-the-world can often be artfully defined to adapt the SEU model to problems that seem, at first, not to fit. The structure as stated might appear to be static and not to incorporate sequential choice possibilities with appropriate adaptive feedback mechanisms, but these complexities are

conceptually accommodated by identifying an action with a dynamic strategy over time.

In the SEU model one assumes that the decision maker can always make up his mind; he has a complete set of transitive preferences for consequences and for lotteries over those consequences. If we keep in mind that consequences can involve a multiplicity of attributes and that action alternatives may incorporate adaptive feedback features, then this assumption of complete comparability – i.e., that the decision maker can always make up his mind in a coherent fashion – is mind boggling. The assumptions imply that the decision maker can accomplish this bit of magic by (a) assessing probabilities for the states-of-the-world, (b) assigning a utility value (real number) to each (complex) consequence, (c) calculating the expected utility value (i.e., the sum of probabilities times utilities) of each lottery associated with each alternative action, and (d) comparing actions by their SEU numbers.

Actually it can be proven as a theorem rather than stated as an axiom that any decision maker who can "consistently" compare actions must behave as if he had embedded within himself a probabilistic assessment over states and a utility function over consequences, and had calculated expected utility values for comparisons. These probabilities and utilities come tumbling out as logical implications of more primitive logical constructs such as full comparability of alternatives, transitivity, and a so-called substitution or sure-thing principle.

Although the SEU model has been treated by many as a descriptive model, it is used primarily as a normative system that captures in a crisp and elegant fashion, the formal properties that characterize one idealized sense of rational choice under uncertainty.

In limited domains the SEU model may also be used as a prescriptive tool in order to guide behavior, but this conscious effort involves a reflective thought process that is far more complex than the bare bones of the SEU model seems to indicate. Real people, in real situations, do not naturally act coherently and one usually cannot discover via their past revealed real behavior their latent probability distributions and utility functions. Rather, the way the SEU model is put to prescriptive use turns the model upside down. We do not start by assuming that the decision maker can, in an unaided fashion, compare any two alternatives but rather we test whether we can compare a few simple hypothetical consequences. Already in this limited domain he might exhibit intransitivities among the few consequences that he is willing initially to compare, but he then must be willing to reflect upon these inconsistencies and modify his preferences so that they line up transitively. In an iterative fashion he must be willing, in a particular instance, to act quite unnaturally: to deliberately police his choices in hypothetical simple situations, one by one, and force them to conform to the desideratum of consistency. Gradually, if he is successful, a probability distribution over states and a utility function over consequences will emerge. These will literally have to be constructed bit by bit, and it is a Platonic myth that latently these probabilities and utilities really exist deep down and that the

analyst merely has to cut away the fat in order to display the pre-existing structure. Next a leap of faith is required: the decision maker must be willing to use his probability and utility functions that he has laboriously constructed to calculate SEUs that will guide his selection of real-world alternatives.

Why should anyone behave in such an unnatural fashion? Well, first of all, in some simple situations for which probabilities are crisply given and based on relative frequency data, and for which consequences can be adequately described by a single numeraire, like money, the SEU prescriptive process is reasonably transparent and natural. Secondly, there is always the nagging question: is there a better alternative? Remember we are assuming that there is a decision maker who is confused about what he should do in his own best interests. (This in no way implies that these interests have to be hedonistically self-serving.) Thirdly, the prescriptive use of the SEU model is appealing to a lot of people who have thought hard about these problems. The same of course, could be said of other prescriptive procedures. Unfortunately, for those of us who are emotionally and intellectually wedded to the SEU prescriptive process, we cannot cite statistical evidence of the form: 100 decision makers were randomly selected; 50 were instructed in the intricacies of the SEU prescriptive process; 50 others in a control group were not. Of the 50 so instructed, 42 became rich and beloved, whereas 7 of the control group were so indicated. It is hard to accumulate descriptive data on the real use of this prescriptive process. Some favorable laboratory data could be cited, but to the skeptics such data are not very convincing.

A decision maker who constructively employs the SEU model to guide his choice must decompose his judgments about uncertainties from his preferences for consequences. After a separate analysis of each component is made, the analyst fuses these elements together to arrive at a choice. But contaminating influences will often permeate across a boundary that is meant to keep the component activities separate and pure. A decision maker's concerns about utilities might influence his assessments of probabilities and vice versa. In addition, before any analysis is done he may have preferences for actions that might influence his assessments of both beliefs and values.

We can be glib in normative theories by hypothesizing the existence of decision agents who can think separately and distinctly about uncertainties and values and who can then integrate these deliberations jointly to determine preferences for actions. However, real behavior often does not conform to such an ideal, especially in complex, highly emotional choice situations. The sophisticated prescriptive intervenor, who wants to help a real client to decide wisely, must be cognizant of these realities. The prescriptive analyst and client must work carefully to ameliorate some of these potential distortions. Sometimes it might be best to give up: to discard as nonimplementable a formal structuring and decomposition of the problem. But let us not go too far and assume that every case presents insurmountable practical difficulties. In prescriptive analyses, especially when inputs are required from diverse experts, each of whose expertise

pertains to only part of the problem, there may not be any other recourse than to decompose the problem into component parts. The point of all this is that the complications alluded to above are real and must be addressed in prescriptive analyses and this is one feature that makes prescriptive analysis so different from normative analysis.

Let us return to an earlier observation that preferences for consequences are not stable, not consistent, and not precise. In descriptive empirical studies these observations are repeatedly documented. Abstract models could try to capture this reality by positing an error theory. A preference value for a consequence could, at any instance, be considered a random drawing from some distribution. Distributions might overlap and therefore at one instance consequence A might be preferred to B and an instance later it might shift. Descriptive modelers have developed a host of other accounts to accommodate the empirical data. In normative analyses one can posit that an accumulation of information or a slight change in one's external environment can account for these shifting preferences even over short periods of time. In prescriptive analysis the spirit is often different than either the normative or the descriptive analysis. A subject may be gently confronted with the observation that some of his responses are seemingly inconsistent and he may be invited to think more deeply about these inconsistencies. Ideally, of course, one might hope that with deeper reflection, with a clarification of the descriptions of consequences, with a better comprehension of the relevant attributes that should be contemplated, the subject would settle down to a set of coherent preferences that could be articulated. Anecdotes can be cited where this state occurs. But, in some cases where preferences shift into a coherent mode, there is the obvious danger that the subject might only be misleading herself and her analyst; she may be generating consistent but not "intrinsically truthful" responses. Anecdotes can also be cited where deeper reflection exacerbates the confusion and coherence is never achieved. It is not generally understood, especially by critics who interpret the SEU model literally, that in a given choice situation it is rarely necessary to have a full articulation of the decision agent's full set of preferences. Occasionally a simple break-even analysis or a sensitivity analysis centered around some crude approximation of the subject's preference structure will suffice. A decision maker may not be sure whether a parameter should be 10, 15, or 20 and in successive probings he may register labile values; but a simple break-even analysis might show that it really does not matter: the break-even point may be 40, way out of the hazy range. At other times an analysis can proceed without obtaining any definitive trade-offs between objectively incommensurable qualities; the answer might be obvious by a simple dominance argument. At other times the prescriptive intervenor might be able to employ an incomplete description of the client's preference structure to eliminate some but not all contending alternatives. In an iterative fashion, the analyst might learn what incremental information should be elicited from the client in order to help the client arrive at a wise choice. This type of intellectual activity is not adequately captured in the

usual descriptive or normative literature. The art and science of elicitation of values (about consequences) and judgments (about uncertainties) lies at the heart of prescriptive endeavors.

As we noted above, individuals, by and large, do not follow the precepts of the SEU model. They do not naturally decompose their choice problems into concerns about beliefs and values; they do not base their choices in real situations on consideration of what they think they would want to do in simple hypothetical situations. These observations are readily acknowledged by the "prescribers" and they respond, "If people behaved naturally as we wish they would behave, there would be no need for our services." The discrepancy between theory and behavior is the *raison d'être* of prescriptive interventions.

In the next section we examine discrepancies between descriptive and normative perceptions in choices where probabilities are crystal clear. Afterwards we grapple with cases where distortions about probabilities are the heart of the matter.

PSYCHOLOGICAL CONSIDERATIONS

Tversky and Kahneman (1986) show that most individuals violate the implications of the SEU model even in very simple choice situations where there is absolutely no vagueness about uncertainties – i.e., where probabilities are identifed with known relative frequencies. They go on to catalog the types of discrepancies between the SEU model and descriptive behavior and then they indicate how the SEU model can be modified to better predict actual behavior. Their descriptive, empirically based, model is called "prospect theory."

Prospect theory is a description of how people make decisions between simple lotteries. Three systematic differences between prospect theory and SEU are demonstrated convincingly by them. First of all, people think of consequences as increments (or decrements) to current wealth and have aversion to losses. Loosely speaking, most individuals exhibit an inflection point at their status quo reference point: they are concave (i.e., risk averse) above it and convex (i.e., risk seeking) below it. Secondly, they do not distinguish adequately between large numbers. Twenty thousand dollars sounds a lot like twenty-five thousand dollars. Thirdly, people give unlikely events more weight than they deserve, and give correspondingly less weight to very likely events.

In decision-tree analysis the usual prescription is to accumulate financial flows along the path of the tree and register their total at the end of the tree before working backwards. If a flow totals $1,000, it is irrelevant if the component summands are $500 and $500 or if they are minus $200 and $1,200. But it is not relevant according to the findings of prospect theory where subjects view sequential choices in terms of their reference positions along the way. That is descriptive reality!

In SEU analysis most utility functions are taken to be concave so that a gamble

is less preferred to its mean. But in prospect theory more subtle distinctions are shown. For desirable payoffs above the reference point there is a strong certainty effect; if subjects are pleased with an outcome that is certain, they hate to gamble it away. But, if the payoffs are undesirable relative to the reference point, subjects have a reversal effect: they become risk seeking.

For those of us who want to help people make better decisions – better for them, not for us – these results pose dilemmas. We can deftly structure problems to avoid such tendencies. For example, in teaching decision analysis at the Harvard Business School, the instructors insist that analyses be made in terms of final financial asset positions rather than in terms of increments or decrements from a present asset position because, in terms of that accounting, the students will fall prey to a psychological trap, to a cognitive illusion. But the problem persists: are we doing right to force students to think one way if they feel another way? That cannot be right. A better way would be to lead subjects to become consciously aware of those behavioral tendencies that contradict the desiderata of coherence. If a subject behaves in a way that makes it possible for someone to make a book against him, then he might choose to live with that behavioral anomaly or to revise his thinking and feeling. How far should we push this indoctrination? These concerns – and they are troublesome – do not fall in the province of descriptive or normative analysis as normally interpreted. They are at the heart of the craft of a discipline that could be called "prescriptive decision analysis."

A deeper question lurks in the background. If a subject mechanically uses the SEU model, many of these so-called behavioral anomalies are eliminated. But is this in the decision maker's best interests? Of course we may construct situations in which decision makers, by following their natural tendencies, are led into obvious errors, but people do not spend their lives entering into deals with devious decision analysts. They deal with the world as it is. People may have learned by experience that unlikely events happen more often than they should, that adverse selection causes alternatives with the potential for loss to produce losses more often than an objective analysis would suggest. Does the SEU model deal with real life decisions, or only with objectively verifiable decisions?

Prospect theory does not answer these questions; it merely poses them. From the perspective of a person wedded to the SEU model, prospect theory confirms the need for such a model to guide behavior. For the "Darwinian School" (that believes 40 million years of evolution have produced humans with an effective capacity for decision making that should not be tampered with) prospect theory confirms the ivory tower nature of SEU.

There certainly are discrepancies between the tenets of SEU theory and descriptive behavior. Should we be changing the analytical procedures we propose for guiding real behavior or should we be applying a bit of psychological therapy? Or both? But before we consider those questions let us first reflect on discrepancies between theory and behavior on the probabilistic (as opposed to the utility) side of the ledger. We will see that on the probabilistic side

we will be more inclined to be more paternalistic: we theorists are not going to change, let the subjects shape up!

BELIEFS OR JUDGMENTS ABOUT UNCERTAINTY

The SEU model asserts that the decision maker should hold beliefs about uncertainties that are in accordance with some subjective probabilistic measure. But copious examples show that most individuals do not "believe" as the SEU theory says they should believe.

Ellsberg (1964) and others have shown that subjectively scaled probability assessments do not always obey the usual probabilistic axioms. Dramatic examples have been concocted for which subjectively assessed probabilities for an event E and for not-E do not sum to unity. Loosely speaking, there is a manifest bias against vagueness, and this discrepancy can lead to choices that are not in conformity with SEU behavior. "That's fine," say the normative analysts, "this shows that we have something to teach."

Tversky and Kahneman (1983) demonstrate empirically that many individuals register beliefs that imply, for them, that the joint event (A and B) is "more probable" than event A alone, the "conjunction fallacy." For example, the probability of a nuclear war starting by a terrorist act is assessed as being more likely than a nuclear war starting. Normative analysts are not disturbed at this: "Oh, that's just a mistake; mention of a terrorist act just helps the imagination." Prescriptive analysts must be aware, however, of these common "mistakes" in the elicitation of judgments about uncertainties.

There are a host of examples where untutored individuals, and even many "tutored" ones as well, have poor intuitions about probability. They might use heuristics developed for one class of problems for another class, such as: information from sample sizes less than 30 should be ignored. Most individuals have poor intuition about how information should modify judgments about uncertainties. Examples abound where even experts neglect base rates, where doctors confuse $P(A/B)$ and $P(B/A)$, where optimism or fatalism or guilt or religious fervor affects deep beliefs about uncertainties. A Peanuts cartoon depicts wishful thinking: "Can this happen?" "Yes, one chance in a million." "Well then let's play." They play; they lose; they become angry for having lost. This is reality.

It is hard to shake most individuals' strong beliefs about the gambler's fallacy: "I've had a couple of successes and therefore I am due for a failure." And there is also the countervailing fallacy: "The dice are running hot. I better get into the act."

The above examples illustrate that often people's beliefs about uncertainty are not in accord with empirical relevant frequencies of plausible physical models – e.g., "a die hath neither memory nor conscience." Normative analysts can remain aloof: "Probability theory does not describe how people think but how they should think." But prescriptive analysts must beware lest they build a so-

called logical superstructure on a nonsensical subjective base. Responsible, professional prescribers must be cognizant of such behavioral anomalies and guide their clients around the common pitfalls.

There is another class of common errors where probabilities get linked with outcomes and alternatives. Let p be the probability of an adverse outcome in the population (e.g., a particular form of cancer). An action is contemplated that will reduce this probability from p to $(p - \Delta p)$ and thus save $N\Delta p$ potential victims (where N is the population size). But administrators and the public might evaluate the efficacy of the proposed action not on Δp but on $\Delta p/p$. A reduction from $p = .18$ to $p = .14$ may be deemed much less desirable than a reduction of $p = .04$ to .02, say. "Oh well," someone might retort, "this might not be such a mistake if one includes public anxiety as a reality of the problem – even if this anxiety is based on misperceptions of the public's own interest." These examples raise the question of what concerns and possible misperceptions should and should not be included in prescriptive analyses.

If the SEU model is to be applied to guide behavior, then someone will have to supply the basic inputs: probabilities and utilities. Let us continue with the probabilistic side. Many, if not most, real decision problems cannot be analyzed adequately using purely objective probabilities. Subjective assessments must be introduced and this once again leads us into a confrontation between abstract theory and realistic behavior. Real people just do not behave like the models say they should, but still they might need and want help. Descriptive theorists happily demonstrate that people are incoherent in their probabilistic assessments; normative analysts generally remain aloof and do not get involved in empirical measurements; it is the prescriptive analysts who must learn how to elicit judgments and make sense out of them, if there is some sense to be gleaned from those judgments.

We have learned that many experts are willing to answer hypothetical questions about uncertain quantities, such as: Do you think it is more likely that X will fall in the interval from a to b or outside that interval? We have learned that lay people and experts alike do not calibrate well: by and large assessed probability distributions are too tight; people think they know more than they really know and are surprised far too often. We know that some assessment methods lead to less distortions than others. We know that it is devilishly hard to assess small probabilities. We know that subjects can learn to calibrate better if they are given systematic feedback. We know that there is a need to develop better methods for elicitation and that the describers and abstracters are not going to provide the impetus for this development. It is the prescribers who must take the lead.

On the probabilistic side we tend to view incoherent responses as errors in perception that should be monitored and corrected. The matters are more complicated on the utility side. Incoherent utility judgments may also result from fallible thinking but there are other possibilities as well: the analysts may be using an inappropriate analytical framework, or the analysts might have abstracted

away too much of reality. We have left the discussion of this problem to the last section because the contrast between probabilities and utilities is instructive.

PSYCHOLOGICAL CONCERNS AS CARRIERS OF VALUE

The study of risky choice began with the study lotteries where all outcomes are simply monetary gains or losses and this paradigm is still the major focus of research in the field. However, the axioms of utility are also applicable for decisions involving nonmonetary and multidimensional outcomes such as health, the environment, or social welfare. In multi-attribute utility theory, the decision maker is asked to identify the dimensions of an outcome that are important determinants of preference. There are few normative guidelines covering what attributes are legitimate and how they should be traded off.

Let us examine the choice of a lottery from the perspective of multi-attribute utility theory. What are some of the concerns of a decision maker faced with a risky choice? A somewhat different question may make the issue clearer: why do many people wish to avoid risk? Bell (1982) has looked at the implications of two *ex post* psychological conditions: regret and disappointment. Regret may occur when a risky choice turns out to be "wrong" after the fact: an alternate choice would have been better given the state-of-the-world that occurred. It is a great frustration to take a chance on a shortcut to the airport only to find that you would have caught the plane had you only kept to your usual route. Not only have you missed your plane but you must live with the fact that it was you, and not providence, who were responsible for the error (see, e.g., Kahneman and Tversky, 1982). Many people anticipate the possible (post-decisional) regret in such situations and trade off a higher chance of making the plane in favor of reducing the possibility of such regret.

Disappointment is a psychological reaction to an outcome that is below expectations. Many people would rather not be told they are being short-listed for an important promotion because this would raise their expectations and lead to great disappointment should they not, in fact, be given the position (see Bell, 1985).

We all recognize that the feelings of anxiety, joy in anticipation, regret, disappointment, elation, envy, and others are a constant part of our lives, but what role should they play in prescriptive analysis? Natural questions for study are (1) To what extent do these psychological concerns actively affect the decisions that are made? (Descriptive.) (2) Should they affect decisions that are made? How can these cognitive concerns be incorporated into more complex models? (Normative.) (3) What are the implications of these cognitive concerns for the way we help people make risky decisions? (Prescriptive.)

It has been argued by many that behavior is always rational; apparent violations of the canons of rationality stem from a too restrictive view of human motivations. Normative modelers might want to embellish the canonical model of willful-choice behavior to include psychological concerns such as envy,

anxiety, joy in suspense, regret, disappointment, elation. Certainly the SEU model can be extended to capture some of these nuances by merely elaborating the description of consequences. But does this capture the essence of the problem? Once research modelers move into this psychological domain they would have to sort out what is fundamental from what is derivative, and they would have to grapple with deep problems of the "divided self."

Economists like to point out that any trade-offs that are made to mitigate or exaggerate psychological satisfaction will diminish real economic benefits. If people behave coherently in an extended psychological sense but not coherently in terms of economic payoffs, then a "book" can be made against them. In social poker, for example, a player might choose not to maximize his expected winnings but rather to bluff more than is empirically profitable, or "stay in" too often merely for the fun of it; he may be maximizing his satisfaction, which may not be congruent with maximizing monetary payoffs. Other players, of course, can exploit these behavioral motivations.

Not all deviations from the restrictive SEU model deserve to be rationalized and made legitimate in an elaborated model. Perhaps an analyst might want to incorporate regret and disappointment but how about errors that are made in subjectively updating probabilities? How about enriching a normative model to accommodate the behavior of individuals who enter into contracts which will purposely restrict their future choices because they now know what they later will want to do is not what they (now know is what they) should do. If a researcher decides to extend the classical normative model to include such concerns, what meta-criteria should be used to decide which embellishments should be included? Can we order heuristics and biases as to "eligibility" or normative status? Are any of these questions empirically answerable?

Or maybe the research agenda should be pitched the other way, not towards the reform of the underlying conceptual model but towards the reform of subjects. Perhaps they need education and/or therapy. In some cases therapy might be achieved by an explicit recognition that some cognitive concerns should be acknowledged and once explicitly analyzed these concerns might turn out not to be so important after all. Take the following example. We have explored Smith's preferences for money and have determined that she should prefer a 50–50 gamble on $10,000 or nothing to the certainty of $2,000. But Smith says, if the chips were down, she would take the $2,000 despite her earlier responses to utility questions. When pressed it appears that she is terribly concerned about the anticipation of the possible post-decision regret she would have if the gamble yielded the $0 payoff. "I would feel terrible ending up with this zero value, knowing that I could have had $2,000." Yes, she undoubtedly would feel terrible, but how terrible? She should also consider that there would also be post-decision delight if she ended up with $10,000 and some ambivalence if she accepted $2,000 knowing that she was giving up a gamble with an expected value of $5,000. In other words, each monetary consequence carries along with it some psychological baggage. If we came to grips with these psychological concerns by trying to cost

them out – e.g., how much would you be willing to pay in dollars to be able to wave a magical wand and get rid of guilt feelings? – the subject might discover that those concerns left unanalyzed were magnified in her mind and once analyzed were not as important as she thought. Such analysis might alter feelings as well as actions. Therapy through analysis.

The reader can hardly fail to notice that we have raised many questions and answered very few. The message is simply that the art and science of decision analysis require many skills not readily classifiable as either normative and descriptive, but which nonetheless have a legitimate and important role in both research and practice.

REFERENCES

Bell, David E. (1982). "Regret in decision making under uncertainty," *Operations Research* 30, 961–81.

(1985). "Disappointment in decision making under uncertainty," *Operations Research* 33, 1–27 (chapter 16, this volume).

Ellsberg, Daniel (1964). "Risk, ambiguity, and the Savage axioms," *Quarterly Journal of Economics* 75, 643–69.

Kahneman, Daniel, and Tversky, Amos (1982). "The psychology of preferences," *Scientific American* 246, 160–73.

Keeney, Ralph, and Raiffa, Howard (1976). *Decisions With Multiple Objectives*. Wiley, New York.

March, James (1978). "Bounded rationality, ambiguity and the engineering of choice," *Bell Journal of Economics*, 9 (2), 587–608 (chapter 2, this volume).

Tversky, Amos, and Kahneman, Daniel (1983). "Extensional v. intuitive reasoning: the conjunction fallacy in probability judgment," *Psychological Review* 90, 293–315.

(1986). "Rational choice and the framing of decisions," *Journal of Business* 59 (4), part 2, 5251–78 (chapter 9, this volume).